Suppliers and Demanders of Flexibility: The Demographics of Gig Work*

M. Keith Chen

UCLA Anderson School of Management

Judith A. Chevalier

Yale School of Management

Peter E. Rossi

UCLA Anderson School of Management

Lindsey Currier

Uber Technologies

Abstract

Platform gig work such as rideshare driving involves workers supplying flexibility to the platform, for example, providing service when demand is high. It also can be attractive to workers who demand flexibility, for example, workers with irregular commitments in other jobs. Who benefits the most (and least) from flexible work arrangements? Workers who supply labor price elastically provide flexibilty to the platform and receive above the platform-average compensation. In contrast, workers with the most time-variation in their reservation wage are demanders of flexibility and benefit from the availability of flexible work options. Using an empirical Bayesian model, we estimate driver-by-driver both the level and time variation in the driver reservation wage. We characterize the demographics of Uber drivers and explore the characteristics of drivers who supply flexibility and the characteristics of drivers who would drop out if the arrangement were less flexible. Our results run counter to several common intuitions about the costs and benefits of gig work.

*Uber Technologies provided the data used in this paper. Under the agreement granted the authors, Uber has the right to review the paper for confidential information but not to dispute or influence the findings or conclusions of the paper. Chevalier and Rossi have no material financial relationships with entities related to this research. Chen is a former employee of Uber, and as a result, continues to hold stock options that may constitute a material financial position. Currier is an Uber employee. This research was found to be IRB exempt (Yale approval case 2000024805).

1 Introduction

At-will labor relationships mediated by digital platforms have become a visible and growing part of the gig economy. Platforms such as Uber, Lyft, TaskRabbit, Door Dash, Bird, Lime, and Instacart all rely on gig workers. Workers on these platforms perform services in response to demand but, critically, have flexibility in deciding whether or when to work. That is, for example, a worker who chooses to work recharging Bird or Lime scooters one night typically has no contractual obligation to charge them the next. Because of this flexibility, many digital gig economy participants use gig work as a supplement to another economic activity such as a primary job, household production, entrepreneurial activities, or education. In this sense, gig workers can act as both suppliers and demanders of flexibility. They supply flexibilty to the platform by working nonstandard hours and working when demand for their services are high. However, some demand or consume flexibility in that they wouldn't work for the platform if the platform didn't allow for at-will working hours. While the demographics of gig workers have been addressed in the literature, it is less well-understood which types of workers supply work very flexibly and which would not be willing to supply labor if these platforms were less flexible.

In this paper, we use data from nearly two hundred thousand drivers on Uber (a popular ride-sharing platform), to examine the demographics of gig work flexibility. Some of these gig workers would be willing to work in an environment in which hours are pre-set and less than fully flexible. These workers value the pay of gig work but not necessarily the flexibility– gig work happens to be a job available to them but they would undertake non-gig work if it were available. Some of these workers disproportionately supply flexibility, providing labor that is very responsive to demand (and pay opportunities) on the site. Some gig workers, due perhaps to the importance of their other commitments, value the flexibility of gig work and would not commit to supply labor to a less flexible job if the gig opportunity did not exist.

To examine the supply of flexibility, we straightforwardly examine the propensity of drivers to supply labor at high payout times. High payout times will tend to be those in which there are a lot of demanders relative to drivers in the marketplace; driver utilization is high and per-mile fees may be elevated. Drivers can provide flexibility by consistently working high earning hours, for example, every Friday after midnight, or by responding to a random shock, such as a large sporting event. We examine the demand for flexibility by examining driver surplus from flexibility. Following the approach in developed in Chen et. al. (2019), we identify the taste for flexibility as being driven by (and equated with) time variation in a worker's reservation wage. If a worker had a constant reservation wage in all hours, the worker would be indifferent between a job that prescribed which specific hours the worker worked and a job that let the worker choose his or her hours, holding all else constant. This time variation in a worker's reservation wage can result from stable differences in the mean reservation wage across time periods, for example, a preference to not work late nights. Time variation can also derive from transitory shocks to reservation wages. For example, a parent may have a very high reservation wage on a day that a child is home sick.

To examine demand for flexibility, we borrow the identification strategy of Chen et. al. (2019). We use data from drivers' decisions of whether and when to supply labor on the Uber platform to estimate each driver's pattern of mean reservation wages for different time blocks and also estimate the variance of each driver's reservation wage due to shocks. This allows us to estimate driver surplus from driving for Uber and to estimate changes to the driver's labor supply and total surplus that would result

from requiring the driver to instead work specific patterns of hours. Using new data provided by Uber Technologies, anonymized driver data is matched to driver demographic characteristics. These data allow us to identify characteristics of drivers that particularly value flexibility and for whom participation in the platform is dependent on that flexibility. It also allows us to identify characteristics of drivers whose participation is particular price-elastic.

Our identification strategy, loosely speaking, is simple: if we see a driver supplying labor in an hour when the expected wage is \$15/hour and choosing not to supply labor in an hour when the expected wage is \$25/hour, controlling for a variety of other factors, we can infer that the driver's reservation wage is time-varying. The pattern of when the driver drives reveals, for each driver, characteristics of the driver's reservation wages. This analysis allows us to estimate whether the driver would supply labor to Uber if the environment were more restrictive about the pattern of hours drivers must drive. It also allows us to estimate the extent to which the driver's labor supply would change with changes in the payouts from driving.

We are interested in both part-time Uber drivers and full-time Uber drivers. As documented by Campbell (2017), most rideshare drivers obtain a minority of their household income from driving, suggesting that driving is often a secondary economic activity. Many part-time drivers who demonstrate high reservation wages during some hours presumably do so due to the time demands and remunerativeness of the the driver's other economic activities. It is unsurprising that some drivers would not drive if it couldn't be worked around the primary economic activity.

Because of this, it may be ex ante difficult to speculate which demographic groups would be most likely to value flexibility and which groups would be most likely to withdraw their labor supply from Uber were it to adopt a less-flexible scheduling regime. For example, while a common intuition is that higher income or wealthier people and women value job flexibility more than men, that is not necessarily true in the Uber environment, where many drivers schedule their work around other, likely less flexible, jobs.

Our paper proceeds as follows. Section 2 briefly reviews the literatures on job flexibility, dual jobholding, and on gender and demographic issues in the gig economy. Section 3 describes our data sources and construction of the analysis dataset. Section 4 provides a first look at the habits of Uber drivers of different demographics. Section 6 briefly reviews the labor supply model introduced in Chen (2019) and outlines how we conduct inferences for that model. Expected labor surplus, labor supply , and expected labor supply are discussed in Section 7. Section 8 provides a conclusion and summary of our findings.

2 Literature

As most rideshare drivers derive a minority of their overall earnings from driving, we consider driving to be closely linked to dual job holding. Substantial research suggests that multiple job holding has historically been limited to about 5 percent of the workforce, although it is more prevalent for workers in certain occupations (for example, Lale (2015) reports that multiple-jobholding rates for teachers are no less than 13 percent). While multiple-jobholding rates are low, a much larger number of workers transition in and out of multiple job holding over the lifecycle (see Paxson and Sicherman (1996), Renna and Oaxaca (2006), and Lale (2015)). Lale (2015) estimates that about 1 percent of full-time single jobholders and 2 percent of part-time single jobholders transition in to multiple jobholding each

month. Uber drivers similarly have high churn. Cook et al. (2018) demonstrate that many workers drive for Uber only for a short time.The literature also suggests persistent geographic differences in dual job holding. Hirsch et. al. (2017) demonstrate that multiple job holding is weakly pro-cyclical, suggesting the importance of labor demand, and that it is negatively correlated with commuting times. This is at least suggestive that technologies that render secondary work more flexible may attract workers to secondary work. The extent to which workers use contingent contract worker is difficult to ascertain from conventional government statistics. The BLS's recent study of contingent and alternative employment relationships only studies workers who report contingent and contract work as their main job. Koustas (2018) uses a large financial aggregator and finds a lower bound of 22% of drivers additionally that have consistent non-rideshare employment through their first quarter of driving.

Clearly, the flexibility of Uber is important for some drivers. Hall and Krueger (2016) examine survey evidence and Uber administrative data. They document that drivers cite flexibility as a reason for working for Uber and that many drivers report that Uber is a part-time activity secondary to more traditional employment. Their findings are consistent with the third party survey in Campbell (2018). Campbell (2018) finds that only about one-third of rideshare drivers report that the majority of their income derives from driving. Prior to the introduction of these platforms, there were clearly fewer opportunities to undertake secondary work that could be accommodated around the schedule of the primary work. Using data from individual bank and credit card accounts, Farrell and Greig (2016) present evidence that is strongly suggestive that workers supply more labor to online platforms such as Uber and Lyft when they receive negative shocks to their earnings in their other sources of employment.

Consistent with this, Hall and Kruger (2016) and Chen et. al. (2019) document that the hours supplied by drivers vary considerably from week to week. Chen et. al. (2019) examine drivers' labor supply in more detail. Because of the flexibility of the platform, a driver can decide whether to supply labor minute by minute, which in turn allows us to infer time patterns of the driver's reservation wage. If there are time periods in which there is on average a substantial disamenity value to driving, supply and demand should lead to an equilibrium of higher expected wages during the undesired hours. Both the typical weekly pattern and shocks to the driver's reservation wage can in principle be extracted. Chen et. al. (2019) examine each driver's labor supply decisions and estimate driver response to alternative scenarios which mimic the effects of traditional employment relationships.

There is a very small literature on driver demographics. Hall and Kruger (2016) document driver demographics from a survey. Cook et al. (2018) show that women Uber drivers earn somewhat lower wages per trip than do men. They demonstrate that this is largely due to women being in a lower position on the experience curve and due to women drivers driving more slowly. Caldwell and Oehlsen (2018) estimate Frisch labor supply elasticities for women vs. men and find that women's hours are more wage-elastic than are men's hours. However, to our knowledge, there are not other papers that analyze the relationship between flexibility, demographics, and the decision to participate in gig work.

3 Data Sources and Construction of Analysis Datasets

Our data are provided by agreement with Uber. We use the same base data as in Chen et. al. (2019). That is, for the period from September 2015 to April 2016, we start with the universe of all UberX

driver-hours for drivers in the twenty cities in the United States with the most UberX trips. For our analysis, we divide time into discrete hours as the unit of observation, 168 hours per week. We define a driver to be active in an hour if she is active for at least 10 minutes within that hour, where active is defined as having a passenger or being en route to a passenger. In other words, drivers are defined as "active" when there are up to 50 minutes of idle time within an hour. We then measure the driver's discrete choice of being active in each of the 168 hour blocks. Thus for each driver, our measure of labor supply consists of a vector of 168 hours/week x 36 weeks of zeroes and ones. This is the variable we refer to when we consider the binary outcome of whether a driver "worked" in a given hour.

We also use the Uber "wage" faced by drivers for each city for each hour. Each driver's total earnings in that hour, divided by minutes worked, times sixty is averaged for all drivers in the UberX sample. Here, minutes worked is defined as the minutes a driver is available on the app, regardless of whether they are on a trip. This may be an overestimate, since drivers are free to do any non-work activities they want during idle time while the app is on. However, because we condition on drivers being active for at least 10 minutes each hour, we remove some cases of drivers ignoring the app, refusing to accept trips, or having the app on in remote locations. While individual driver's wages will vary from the city average, we treat the city average as what the driver can expect to make if she or he chooses to drive in a particular hour. It is important to remember that, on the Uber platform, drivers are expected to pay for both the capital costs of their vehicle and all costs of operating the vehicle. In our analysis, these costs are incorporated into the driver's reservation wages.

Because we will be evaluating patterns of activities over time, our analysis sample of drivers consists of drivers who are active in at least 1 hour for at least 16 of the 36 weeks that we have available in our data. We will refer to drivers who meet this criteria as "active drivers." This is an important filter. Our understanding is that many drivers try out driving for Uber but abandon the platform. Cook et al. (2018) find that 68 percent of Uber drivers who start driving for Uber have abandoned the platform after 6 months (though, because Uber drivers do not have to formally quit, it is possible that some are on an extended break). The platform was growing rapidly during the time of our data; our data requirements force us to oversample drivers who remain on the platform for a relatively long time. As mentioned above, we also restricted attention to the top 20 US cities by volume of labor supplied on the UberX platform. This gives us driver data for 196,198 drivers.

Uber internally, of course, maintains data on the identity of their drivers. While Uber did not share driver identities with the researchers, they provided anonymized information on driver demographics for our study. Using this information, we have information on driver age, driver gender, and estimates of driver ethnicity and driver neighborhood income. Driver age and gender come directly from Uber's internal records. Driver neighborhood income is estimated by matching the driver's address to the Census Bureau's geocode resource. This method allows matching to Census tracts for a total of 91% of drivers (mismatches may be due to street misspellings). The median income and the gini coefficient for each driver Census tract was extracted using the 2013-2017 ACS summary file. Because of concerns from Uber that the exact median income figure could compromise driver privacy, the income and gini coefficients were further aggregated. Specifically, the Census tract median incomes for all drivers were put into percentiles and for each driver, the Census tract median income passed to us from Uber is the median for that percentile. The same procedure was used for gini coefficients.

Uber does not collect data on driver race or ethnicity. In order to impute driver race and ethnicity, a procedure similar to that reported in Diamond et. al. (2019) was undertaken. Specifically, for each driver, the racial/ethnic composition of the driver's census block is extracted from the 2010 Census.

Category	Men	Women	% of row Women
Age<=60	135881	25005	15.5%
Age>60	15475	2062	11.6%
White	76817	12703	14.2%
Hispanic	30951	4920	13.7%
Black	24988	8343	25.0%
Asian	18600	1101	5.6%

Tab. 1: Number of drivers in our dataset of 178423 drivers with various characteristics.

Category	Mean Income	% Women
Income Quartile 1	19959	17.2%
Income Quartile 2	27744	15.9%
Income Quartile 3	35851	14.5%
Income Quartile 4	52245	12.9%

Tab. 2: Number of drivers in our dataset of 178423 drivers with various characteristics.

Using driver full names, driver names were used to predict ethnicity using the python package ethnicolr. The package ethnicolr is trained using the method described by Sood and Laohaprapanon (2018) using Florida Voting Registration data. The driver race/ethnicity that we use in our summary statistics is the categorization of white, black, hispanic, and asian with the highest Bayes' rule posterior given the driver's name and census block. Due to the inability to match all driver addresses to Census files, plus missing gender data and age data for a small number of drivers, we are left with 178,423 drivers with complete data.

There are many ways in which this is an imperfect measure of race and ethnicity. For example, this methodology has no mechanism to capture multi-race individuals and treats Hispanic origin as mutually exclusive from black or white. The category for Asian drivers is broad, including East Asian and South Asian and Middle-Eastern drivers. Finally, demographic inferences are an understably sensitive issue (Andriotis and Ensign, 2015). Nonetheless, these techniques have been widely used to better understand ethnic inequalities by government agencies like the Department of Justice's Civil Rights Division, the Consumer Finance Protection Bureau, and the Office of Minority Health in the Department of Health and Human Services (CFPB (2014), Martino et al, 2012). While demographic inference is imperfect, it is in some cases the best alternative, and can be used to increase understanding of important social issues. Finally, note that no method, including self-reporting, is perfect (Arday et al, 2000).

We briefly summarize the demographic characteristics of our driver sample.

Table 1 and Table 2 show some characteristics of our "dedicated driver" sample. Our overall female driver percentage of 15.1% is much lower than the 27% found by Cook et. al. (2019). However, their data runs into a later time period than does ours and they find that the female drivers in their sample have accrued substantially less experience than the male drivers. They also show that women have a substantially higher attrition rate in the first 6 months of driving and our methodology requires

	Share of drivers averaging N hours/week			
	Age<=60	Age<=60	Age<=60	Age>60
Total hours (N)	Female	Male	All	All
1-4	2%	1%	1%	1%
5-12	41%	28%	30%	221
13-20	32%	28%	29%	27%
21-30	17%	22%	21%	54%
31-40	6%	12%	11%	15%
41+	2%	8%	7%	11%
Total drivers in category	25005	135881	160886	17537

Tab. 3: Distribution of average active hours using only weeks in which the driver works at least one hour

driving at least 16 of 36 weeks. In contrast, the New York City Taxicab Fact Book (2014) reports that 1.1% of New York cab drivers are female.

Based on these measures, Uber drivers are surprisingly representative of the demographics of the Metropolitan Statistical Areas of the 20 Uber cities that we study. The driver sample is somewhat less likely to be Hispanic (22% of drivers) than the general population of the MSA (29% of people), with the other racial and ethnic groups correspondingly oversampled in the driver population.

4 Model-free Evidence on Driver Demand for Flexibility for Demographic Groups

4.1 Uber Driver Labor Supply by Demographic Group

Chen et. al. (2019) demonstrate the tremendous variation in Uber driver hours as well as the volatility across weeks in driving behavior for individual drivers. Here, we summarize differences across various demographic groups in driving behavior.

Table 3 examines weeks in which our sample drivers drive. We examine the distribution of average hours worked by various driving types.

Recall that we consider a driver active in any hour when she was active for at least 10 minutes, and we count how many of the 168 hour blocks in the week the driver was active. Figure 3 displays the percentage of hours averaged across those weeks that each driver was active. First, we see that, for prime age drivers, women tend to work considerably fewer hours than men. Second, we see that drivers over the age of 60 work more hours in weeks they work than drivers under the age of 60. This would be surprising in an environment in which most drivers were using Uber as a primary source of income but unsurprising given the extent to which Uber is used as a secondary economic activity, perhaps especially among prime-age drivers.

We also examine differences in driving habits for drivers by race and income.

Table 4shows very similar patterns of driving across racial groups, with the exception that a higher percentage of Asian drivers appear to use Uber as a more full-time activity than any other group. In unreported results, we also examine driving by quintile of census tract income. The propensity to drive many versus few hours is virtually identical across the five income quintiles.

	Share of drivers averaging N hours/week			
Total hours (N)	Age<=60 White	Age<=60 Black	Age<=60 Asian	Age<=60 Hispanic
1-4	1%	1%	1%	1%
5-12	32%	31%	23%	30%
13-20	29%	30%	25%	30%
21-30	21%	20%	23%	22%
31-40	10%	10%	15%	11%
41+	7%	7%	12%	6%
Total drivers in category	73934	30038	20134	36771

Tab. 4: Distribution of average active hours using only weeks in which the driver works at least one hour

		Male			Female		
		% working that block in week t+1			% working that block in week t+1		
Did a driver work	Yes	52.9	55.3	67.9	45.1	48.0	64.1
a block in week t?	No	10.0	12.4	21.2	8.5	10.9	21.0
Conditional			in week	that day		in week	that day
on working		unconditional	t+1	in wk t+1	unconditional	t+1	in wk t+1

Tab. 5: Conditional Probabilities of Working a Block in Consecutive Weeks

4.2 Within-Driver Variation in Schedules

We examine the extent to which drivers of various demographics vary their schedules from week to week. Intuitively, the variation in driver schedules (combined with available wages and the economic incentives to work particular hours) will motivate our model of driver behavior.

To summarize the data, we divide the 168 hours of the week into 56 three-hour blocks ordered sequentially from the beginning of the week. We examine the question: if a driver drives in a block in week t, what is the probability that the driver drives in that same block in week $t+1$? Then, to provide insight into the ways that a driver can alter her schedule, we ask the same question, but condition on the driver working at some point in week $t+1$. The idea is to identify the extent to which week-to-week variability is due to sitting out the entire week. Next, we trace working in the same block across weeks, but condition on driving sometime in the relevant day. The results for men vs. women are shown in Table 5.

Table 5 shows that a male driver who works in a particular block has a roughly 53 percent chance of working in that same block on the following week. If the driver did not work in a particular time block in week t, he has only a 10 percent chance of working in it the following week. The probability that a driver who worked in a block in week t will work in it again in week t+1 increases very little when excluding drivers who take the entire next week off. However, conditional on working sometime that day in the next week, the probability that a driver works in the same three-hour block that he or she worked in the prior week rises to about two-thirds. This suggests that the particular hours driven by a given driver vary considerably, even conditioning on the driver working sometime in the day. This is true for both men and women, with the probability that a woman works the same block in the adjacent week being lower. Our model will allow us to parse this out more clearly, but this is suggestive that women have somewhat less predictable schedules than do men.

		Age > 60			Age <= 60		
		% working that block in week t+1			% working that block in week t+1		
Did a driver work	Yes	58.0	60.1	72.8	51.4	53.8	66.9
a block in week t?	No	10.2	12.0	19.7	9.8	12.2	21.4
Conditional			in week	that day		in week	that day
on working		unconditional	t+1	in wk t+1	unconditional	t+1	in wk t+1

Tab. 6: Conditional Probabilities of Working a Block in Consecutive Weeks

Table 6 shows this same transition comparison for older versus younger drivers. Here, we can see real differences in the propensity of drivers to drive at the same time from week to week. For example, drivers over age 60 who are working in a particular day are less likely to work in a block they didn't work in prior week than are younger drivers, but substantially more likely to work in a block that they did work in the week past.

This evidence of the volatility in driver hours presented above fundamentally does not allow us to disentangle two sources of week-to-week variation in hours worked for specific drivers. Drivers who drive less predictably could be supplying flexibility to Uber– that is, driving when expected payouts are high. However, they could also be taking advantage of the flexibility of Uber– driving when it is convenient for them. We pursue a model to allow us to disentangle these factors. Drivers of different demographics may be differentially unpredictable either because they systematically differ in their propensity to chase high wage opportunities or because their reservation wages are more volatile.

5 Suppliers of Flexibility

Drivers contribute to the Uber platform by providing labor. An additional driver in a local area is particularly valuable to riders on the platform when wait times for rides are high and/or when prices on the platform are high (when there is surge). In contrast, if at a particular time in a particular city, many drivers are waiting for riders, an additional driver is not particularly valuable to the consumers. The payout per hour earned on average by all drivers in a particular city-hour is a proxy for the value of an incremental driver in that hour because driver payouts are high when drivers are a high utilization (defined as active time divided by time a driver is on the app) and when there is surge. We have already seen that driving patterns across drivers of different demographics tend to differ. This raises the question of whether drivers differ systematically in the value they provide to the platform by driving during high-value versus low-value period. A driver that tends to drive when payouts are high is, in effect, supplying valued flexibility to the Uber system.

We create several measures to capture the supply of flexibility. Consider a driver i who works in $N_{i,l}$ hour blocks in a particular week, l. Those hours might be ones in which the driver is particularly valuable to the system or they may be ones that are convenient for the driver but not particularly valuable for the system. We first calculate the payout that the driver would have earned if the driver earned the city average wage for the full hour for each of the specific hours driven. Thus, we are eliminating the effect of the driver being particularly lucky or unlucky (or skilled or unskilled) relative to the city average. Thus, for each week, l, day d, and hour, h, we consider the labor supply decision, $Y_{i,ldh}$, for driver i as well as the average prevailing wage, w_{ldh}, where $Y_{i,ldh} = 1$ if driver i is observed to work in hour ldh and 0 if not. Thus, we calculate the total wage $TW_{i,l}$:

Group	% Possible AvgShare	% Expected Possible AvgShareExp	%Idiosyncratic Possible AvgShareIdio
All	73.3%	77.8%	42.1%
Age<=60	73.4%	78.0%	42.1%
Age>60	72.2%	76.7%	42.3%
<=60, female	71.0%	75.8%	41.2%
<=60, male	73.9%	78.4%	42.2%

Tab. 7: Fraction of theoretical max wage by group

$$TW_{i,l} = \sum_{d=1}^{7} \sum_{h=1}^{24} [w_{ldh}] \times [Y_{i,ldh}] \tag{1}$$

and total hour blocks worked $N_{i,l}$

$$N_{i,l} = \sum_{d=1}^{7} \sum_{h=1}^{24} [Y_{i,ldh}] \tag{2}$$

We also calculation the maximum wage the driver could have earned if the driver worked the $N_{i,l}$ highest wage blocks in week, l. That is, we calculate potential wage or $PW_{i,l}$:

$$PW_{i,l} = max(TW_{i,ldh}|N = N_{i,l}) \tag{3}$$

Our overall measure of the propensity of the driver to supply flexibility to the system, $AvgShare_i$ is the average share of potential wages earned by the driver, averaged across all 36 weeks in our sample.

We can further decompose the share of potential wages earned by the driver into two components, the propensity of the driver to drive in the hours in which an additional driver is typically very valuable, and the propensity of the driver to drive in hours in which an additional driver is idiosyncratically very valuable. More specifically, there are hours, such as 5 a.m. on a weekday which are typically very lucrative (due, presumably, to airport trips). There are other hours, such as 3 p.m. on a Saturday which are not typically lucrative, but could be in a given week due to a sports event or concert. Uber tends to inform drivers of upcoming potential busy times and drivers make varying investments themselves in learning about these opportunities. We might hypothesize that drivers in different demographic groups may be systematically different in their willingness to pay attention to and respond to these opportunities. We are interested in, and measure here, both the propensity to drive in a "regular" lucrative hours versus idiosyncratically lucrative hours.

To examine this, we construct a variant of $AvgShare_i$, $AvgShareExp_i$ in which $w_{ldh}$is replaced with the mean of the wages for the particular hour in the city over the 36 week sample, $\overline{w_{ldh}}$ and another variant, $AvgShareIdio_i$ in which in which $w_{ldh}$is replaced with the deviation from the overall mean of the wages for the particular hour in the city over the 36 week sample, $w_{ldh} - \overline{w_{ldh}}$. For this measure, we rescale the measure to be between zero and one.

We summarize these measures for various driver demographics.

Group	% Possible AvgShare	% Expected Possible AvgShareExp	%Idiosyncratic Possible AvgShareIdio
M,<=60,Inc Quint 1	73.9%	78.1%	82.1%
M,<=60,Inc Quint 2	74.0%	78.5%	71.2%
M,<=60,Inc Quint 3	74.0%	78.6%	48.4%
M,<=60,Inc Quint 4	73.7%	78.4%	53.2%
M,<=60,Inc Quint 5	73.8%	78.3%	60.5%

Tab. 8: Fraction of theoretical max wage by group

The results show interesting demographic distinctions. For example, women earn a lower share of the possible wages than do men, both because the work less-lucrative regular hours but also because they don't work. Older workers have a lower share of possible wages and this is driven entirely by their propensity to drive less lucrative regualr hours. The differences across income quintiles are particularly large. The lowest and second lowest income quintiles do not attain substantially different overall shares of the possibe wages, but lower income workers are particularly likely to obtain high "idiosyncratic" wages. We interpret this as suggesting that workers from lower income deciles are particularly like to respond to opportunities to earn above average wages such as concerts or sporting events. In this sense, workers from lower income deciles are particularly important in supplying flexibility to the system.

6 A Review of the Model of the Driving Decision and Inference Procedures

Here, we briefly review the model in Chen et. al. 2019, from which our results are derived. A simple model of labor supply specifies that drivers will supply labor if their reservation wages are less than the prevailing expected wage. That is, for a given period of time (which we take as one hour), we observe the labor supply decision, Y_{it}, as well as the expected prevailing wage, w_{it}, where $Y_{it} = 1$ if driver i is observed to work in hour t and 0 if not. We define "working" in a given hour \as having at least 10 minutes of "active" time engaged in picking up a rider or on a trip. Expected wages are computed assuming drivers are rational and have access to the distribution of wages in a particular city and time. We estimate expected wages by computing the average wage over all Uber drivers in that city and time (see Section 3 above for details).

It should be noted that our measure of prevailing wages is not net of the variable costs of operating a vehicle. Therefore, our reservation wages should be interpreted as a gross quantity as well. Note that if a given driver has a car that is cheaper or more expensive to operate than the mean driver, this difference in expenses would be reflected in the driver's mean reservation wage. Of course, the driving decision is based on the difference between prevailing and reservation wages which does not depend on assumptions regarding the incorporation of operating costs.

6.1 A Model of Reservation wages and driving decisions

We start by providing a simple intuition of our identification strategy (taken from Chen et. al. (2019)). Consider a weekly one-hour period, say, Tuesday, 2 to 3 p.m. For concreteness, assume that the mean

prevailing wage for that hour is $20 in a particular city, and consider a driver who works that hour most weeks. Our estimation would infer that the driver has a mean reservation wage for that hour that is less than $20. Now, suppose that there are some slow weeks where the prevailing wage is around $15 for that hour. If the driver drives most of those weeks too, that suggests that the driver has a mean reservation wage for the hour that is less than $15, and thus, on the more typical $20 weeks, she is getting at least $5 in surplus. In contrast, if the driver does drive the $20 weeks usually, but *doesn't* drive in the $15 weeks, then our estimate of the mean reservation wage of the driver for that hour will be bounded between $15 and $20. This illustrates how the variation in the wage across weeks helps us to pinpoint the driver's reservation wage. With a lot of data, we'd be able to see the wage at which the driver "drops out" from working in the hour. For the driver who usually drives the Tuesday 2 to 3 p.m. hour when the prevailing wage is $20, if the driver doesn't drive that hour in some of the $20 weeks, given her other behavior, her not driving will have to be ascribed to some kind of shock. The extent to which it is attributed to a shock to her hour or day or week will largely be a function of whether the rest of her day/week are also outliers relative to her other behavior. The variance of the shocks experienced by the driver will be determined in part, loosely, by whether we sometimes observe the driver to not drive in that hour when it is *more* lucrative than a typical $20 hour.

We now turn to a more specific description of our methods. The specification of the reservation wage process is crucial to determining the extent to which drivers are able to exercise flexibility in labor supply. As we have documented in section 4, Uber drivers have both predictable and unpredictable patterns of labor supply. For these reasons, we postulate a model of reservation wages with both a predictable mean component as well as a random component that is unobserved by the econometrician but revealed to the drivers.

$$w_{it}^{*} = \mu_i(t) + \varepsilon_{it} \tag{4}$$

Here w_{it}^{*} is the reservation wage of driver i in time t, $\mu_i(t)$ is the mean reservation wage at time t, and ε_{it} is a random shock to the reservation wage that will be resolved, for Uber drivers, before time t. That is, we assume that by at least the beginning of each time period (hour) each Uber driver has realized the shock and therefore simply compares his or her reservation wage for the hour to the expected wage to make a labor supply decision.

While the reservation wage w_{it}^{*} is unobservable to the econometrician, both driver labor supply, y_{it}, and the expected wage, w_{it}, are observed. Driver labor supply, y_{it}, takes the value of one in any hour in which the driver works and zero in any hour in which the driver does not work. In an hour when the driver works, we can infer that the reservation wage is exceeded by the expected wage. Note that the expected wage in a given period can incorporate common knowledge by drivers about predictable events (such as concerts, conventions, and sporting events) that create peaks in demand for Uber services.

Mean Function

The mean portion of the reservation wage process drives the predictable portion of labor supply. For example, if a driver has a regular weekday job, the model can accommodate this with high reservation wages during the 9-5 hours of each weekday. Since these patterns of labor supply vary widely across

drivers, we must provide mean function parameters that vary at the driver level. Even though we have a relatively large number of driver-hour observations, the censoring mechanism applied to the reservation process means that the information content of even thousands of observations is limited. We use a parsimonious specification by 1) grouping hours into blocks associated with a common shift in the mean reservation wage and 2) assuming driver preferences are stable and not allowing for trends or other time shifts. This implies that our mean function is a function only of the day and hour corresponding to time interval t, $\mu_i(t) = \mu_i(d, h)$.

Our mean specification allows for 9 parameters corresponding to the following blocks of hours.

1. MF_am: Monday-Friday, 7 a.m. - 12 noon
2. MF_afternoon: Monday-Friday, 1 - 4 p.m.
3. MF_rush_hour: Monday-Friday, 5 - 8 p.m.
4. MTh_evening: Monday-Thursday, 9 p.m. - 12 a.m.
5. MTh_late_night: Monday-Thursday, 12 - 3 a.m.
6. FS_evening: Friday-Saturday, 9 p.m. - 12 a.m.
7. FS_late_night: Friday-Saturday, 12 - 3 a.m.
8. MSu_don[1]: Monday-Sunday, 4 a.m. - 6 a.m.
9. Base: all remaining hours in the week[2]

Error Components

We have observed that labor supply behavior of Uber drivers has an unpredictable component at the weekly, daily, and hourly frequencies. To accommodate these patterns of behavior, we employ a three-part variance components model for the shock to reservation wages.

$$\varepsilon_t = v_w + v_d + v_h \tag{5}$$

In this model, each of the error components is iid normal[3] over its respective frequency with standard deviations, $\sigma_w, \sigma_d, \sigma_h$ respectively. "w" denotes weekly, "d" denotes daily, and "h" denotes hourly.

[1] Dead-of-night.

[2] Note that each hour block extends from the first minute of the first hour in the block to the last minute of the second hour in the block specification; for example, the MF_am block extends from 7:00 a.m. until 12:59 p.m.

[3] Normal error components imply that the reservation wage process is multivariate normal over the 168 hours that comprise one week. The assumption of normality allows us to specify a model in which the mean of reservation wages can be determined independently of the size or variability of the shocks or unpredictable component of reservation wages. One possible drawback to this assumption is that there is some probability that reservation wage realizations will be negative (this may be very small). Some might suggest modeling the log-reservation wages. While this certainly removes the possibility of negative reservation wages, this assumption creates other undesirable problems. If we assume log-normal reservation wages, then high mean reservation wages are also associated with high variances. This means that we cannot independently vary the degree to which drivers have unpredictable (large shock) patterns versus when they work on average. To take the example of someone with a high reservation wage during the day (due to another work opportunity), the log-normal model would also require that they be more unpredictable during the day then on weekends and evenings. We do not want to impose this sort of restriction on driver behavior.

Thus, each time period (an hour) sees a new realization of the hour shock, v_h, each day a new day shock, and each week a new week shock.

Since each day within a week shares the common week shock and each hour within a day shares a common day shock, this creates the well-known variance components covariance structure that can exhibit very high correlation between periods within each broader timeframe. The error covariance matrix of the reservation wage shock in (5) is block diagonal across weeks.

Our focus will be on using our model to measure driver surplus both in the actual Uber labor arrangement as well as in various alternative scenarios. Uber driver surplus can derive from a variety of factors. First, some drivers will have low reservation wages overall and will derive surplus from the difference between those reservation wages and the prevailing hourly wage. For an extreme case, consider the lonely driver who enjoys driving and talking to customers. This driver is clearly not the marginal driver who sets the wage, and this inframarginal driver clearly earns surplus. Second, some drivers will have reservation wages that are systematically heterogeneous across the hour blocks, and the Uber structure allows the driver to drive only in the lower reservation wage hours. For example, a driver who always works a valuable noon to 8 p.m. job can systematically not work in those hours. This driver earns surplus by avoiding work in those hours but working in other hours when the primary job is unavailable. Third, some drivers will have significant variance in their reservations wages that differ from week to week and the Uber arrangement allows the driver to shift driving hours. For example, an actor can choose not to drive whenever he is called for an audition. A retail worker can work when a shift has been cancelled. Finally, a driver who has fairly uniform reservation wages across hours can earn surplus by "picking off" the high wage hours (such as the bar closing hours or when a concert or other event is taking place).

Details of our estimation methodology are provided in Chen et. al. (2019).

Parameter results

Our results suggest that Uber drivers do not have homogeneous preferences for time of day and day of week. Figure 1 provides scatterplots of normalized mean reservation wage estimates. Recall that each driver has a separate, and possibly, unique mean reservation wage for all of the nine hour-blocks. For example, the four graphs of Figure 1 shows reservations wages for two time blocks for male vs female drivers and drivers over versus under 60 years old. The y-axis in each shows the mean reservation wage for the Friday-Saturday late night hour block relative to the base period (truncated to remove outliers). The x-axis shows the same measure for the Monday-Friday afternoon hour block. For bot hour blocks in all graphs, the mean reservation wages range from a large positive to large negative deviation from the base period estimates, suggesting that reservation wages for these hour blocks, even within driver demographic type, are very heterogeneous. In addition, for all groups, there is a negative correlation between preferences for the Monday-Friday Afternoon (horizontal axis) and the late night block. Drivers who like to drive one tend not to drive the other. Differences among the demographic groups are also apparent. The over-60 drivers tend to have lower reservation wages for the Monday-Friday afternoon time (graphs look shifted to the left relative to the younger group). While there are many women who have below-baseline reservation wages for the late night time periods, they are a smaller fraction of all women than of all men. Overall, 49 percent of prime age men have a preference for late night weekend driving relative to the baseline (have a negative coefficient for Fri-Sat late night) but only 22 percent of older women.

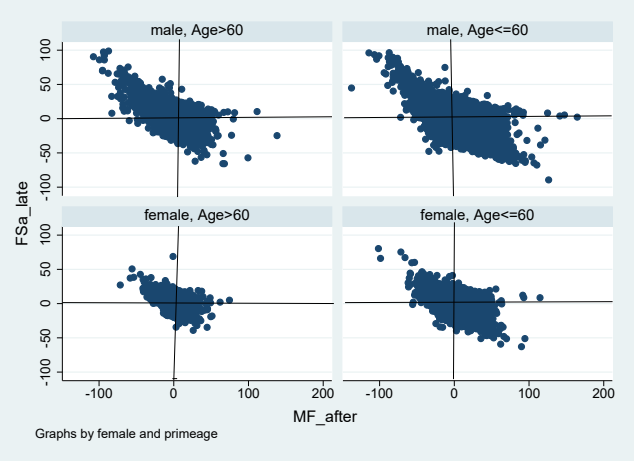


Fig. 1: Scatterplots of Mean Reservation Wage Parameters

The parameter results estimated here play a role both in our estimates of driver's supply of flexibility and the driver's demand for flexibility from the Uber platform. Drivers supply flexibility if they drive a times when wages are high. We will consider both hours when wages are typically high (such as late night) but also times when wages are idiosyncratically high (when there are concerts or other events). Drivers who have a large distaste for particular time are not willing to drive that time even when it is lucrative.

7 Driver Surplus and Labor Supply

The parameters of our model allow us to calculate driver surpus as the difference between the reservation wage and the expected wage in any hour in which the driver is driving. We also can also calculate this surplus for alternative arrangements that afford less flexibility and calculate whether or not a driver would be expected to drive a given hour, as the driver should drive only if the surplus from doing so is positive.

7.1 Surplus Measure

As in Chen et. al. (2019), our goal is to compute the expected surplus for each driver. In our model, drivers will work only if their surplus (excess of wage over reservation wage) is positive. We will compute the expected surplus which is the probability that the surplus is positive (i.e. the driver decides to work) times the expected surplus conditional on working. Consider hour t in which a driver faces wage w_t, expected surplus can written as

$$ES_{i,t} = \left[w_t - E\left[w^*_{i,t} | w^*_{i,t} < w_t\right]\right] \times Pr\left[w^*_{i,t} < w_t\right] \tag{6}$$

To produce a welfare measure for each driver, we sum expected surplus to the driver-week level and compute the average of this measure over all weeks for which we observe the driver in our data. This averages the measure over the distribution of prevailing wages faced by each driver. In the end, we

will have one expected surplus value for each driver. We can gauge the impact of various flexibility restrictions on drivver labor supply and the distribution of this surplus across drivers.

7.2 Constraints on Flexibility

We start with the base case, in which the Uber system imposes no constraints on labor supply flexibility. We will compare the expected surplus under this flexible system with the surplus under an alternative system in which the driver cannot adjust hours in response to the driver's hourly or daily shocks. We can think of the difference between these two surpluses as (some) of the value that the driver gets from the Uber system which informs our understanding of the driver's demand for flexibility. Thus, we define our scenarios as:

(Base) Drivers can adapt to weekly, daily, and hourly shocks with full knowledge of the prevailing wages for that city, week, day and hour and full knowledge of the realization of all of the shocks.

In the base case, drivers make labor supply decisions with full knowledge of the realized value of all weekly, daily and hourly shocks. We consider two other scenarios of decreasing flexibility.

(A) *Cannot adapt to hourly shocks.* In this scenario, we do not allow the driver to adapt to hourly shocks. One interpretation is the driver must make a decision about which hours she'll work at the beginning of each day with knowledge of the distribution of hourly shocks to the reservation wage but without knowledge of the realization of the shocks for each hour in that day. This case affords flexibility to adapt to weekly and daily shocks but not to hourly shocks.

(B) *Cannot adapt to daily and hourly shocks.* Here, we do not allow the driver to adapt to daily or hourly shocks. The driver can adapt to changes in shocks from week to week but not within the week.

It should be emphasized that these scenarios are restrictions only on the driver's ability to adapt to shocks. We still allow the driver to respond to changes in the prevailing wage, and we assume that drivers have perfect foresight as to the prevailing wage. We will examine the driver's sensitivity to the prevailing wage. We also still allow the driver to have a driver-specific profile of mean reservation wages that can vary by day of week and hour of day. That is (A) and (B) are still much more flexible than most conventional work arrangements.

Details of methods for calculating the expected surplus under each of these scenarios is given in Chen et. al. (2019)

7.3 Expected Surplus and Labor Supply Computations

For each of the drivers, we compute Bayes estimates of the mean reservation wage parameter and Bayes estimates of each of the variance components necessary for the expected labor supply and expected surplus computations.

We start with some summary statistics on surplus for different demographic groups. The goal of this exercise is to estimate what demographic groups are those that take advantage of and value the flexibility of the platform.

Surplus		Expected Surplus		
Group	Actual	Base	A	B
All	Surplus	202	84	41
	% of Base		41.6%	20.3%
Age<=60	Surplus	199	82	40
	% of Base		41.1%	20.0%
Age>60	Surplus	230	106	55
	% of Base		46.0%	23.9%
<=60, female	Surplus	146.7	53	20
	% of Base		36.3%	13.9%
<=60, male	Surplus	208	87	43
	% of Base		41.8%	20.8%

Tab. 9: Surplus under alternative scenarios by group

Surplus		Expected Surplus		
Group	Actual	Base	A	B
M,<=60,Black	Surplus	217	86	41
	% of Base		39.5%	18.8%
M, <=60,White	Surplus	202	84	41
	% of Base		41.8%	20.5%
M, <=60,Asian	Surplus	235	106	60
	% of Base		45.1%	25.6%
M,<=60,Hisp	Surplus	200	82	39
	% of Base		41.2%	19.6%

Tab. 10: Surplus under alternative scenarios by group

Surplus		Expected Surplus		
Group	Actual	Base	A	B
M,<=60,Inc Qunt 1	Surplus	210	84	41
	% Base		40%	19.5%
M,<=60,Inc Qunt 2	Surplus	212	88	43
	% Base		41.5%	20.3%
M,<=60,Inc Qunt 3	Surplus	207	86	43
	% Base		41.5%	20.7%
M,<=60,Inc Qunt 4	Surplus	208	88	44
	% Base		42.3%	21.2%
M,<=60, Inc Qunt 5	Surplus	206	88	44
	% Base		42.7%	21.4%

Tab. 11: Surplus under alternative scenarios by group

The tables illustrate that, for all drivers, surplus declines precipitously in scenarios in which we disallow the driver to increase or decrease their labor supply in response to idiosyncratic shocks of the hourly or weekly frequency. The loss of surplus appears to be somewhat larger for women relative to men, for younger people relative to older people, and for lower income groups relative to higher ones.

8 Conclusions

The Uber driver arrangement attracted more than a million drivers to offer labor supply during the 8 month period of our data, which is limited to only the U.S. UberX service. One of the attractions of Uber is the flexibility afforded to drivers. Not only can drivers choose to supply relatively small numbers of hours per week, but they can also allocate these hours flexibly over the days and hours of the week. However, this is not the only or even most important source of flexibility provided to Uber drivers. Another important source of flexibility is the ability of an Uber driver to adapt on an hour-by-hour basis to changes in demands on her time. While traditional workplaces do compete to provide flexibility to workers, the literature suggests that lower-wage, lower-skill workers typically have limited ability to respond to everyday shocks. The goal of this paper is to estimate a model of labor supply, which will allow for a quantification of the value of both flexibility and adaptability.

We postulate a model in which each driver has a reservation wage process with both a predictable mean component as well as weekly, daily, and hourly shocks. This operationalizes the view that workers face unpredictable events that can change their labor supply decisions on an hourly basis. We assume that drivers form rational expectations regarding the expected wage and make labor supply decisions on an hour-by-hour basis by comparing their own reservation wage to the prevailing expected wage. Our model is a multivariate Probit model with a time-varying censoring point which facilitates a greater degree of identification than the traditional Probit structure. Driver-level exact finite sample inference is possible using a hybrid MCMC approach.

We estimate very large shock variances, suggesting the potential for large driver surplus in Uber-like arrangements that allow drivers to decide, on an hour-by-hour basis, when to work. We compute driver labor surplus—accounting for 40% of total expected earnings, or $150 per week on average—under the existing Uber arrangement. Labor surplus for alternative work arrangements, which limit drivers' ability to adapt to hourly and daily reservation wage shocks, are also computed. Constraints on the ability to adapt to shocks have large effects on expected labor surplus; eliminating this ability reduces labor surplus by more than two-thirds.

In summary, we document an important source of value in flexible work arrangements—the ability to adapt work schedules to time-varying reservation wages. Perhaps not surprisingly, this adaptability has high value to individuals who have selected into the Uber platform. Our expectation is that technology will enable the growth of more Uber-style work arrangements. While such arrangements may have important downsides relative to the traditional careers they supplant or supplement, we expect that flexibility will be an important source of value in such arrangements.

References

Harry Campbell. The rideshare guy 2018 reader survey. Technical report, Therideshareguy.com, 2018.

Cody Cook, Rebecca Diamond, Jonathan Hall, John List, and Paul Oyer. The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. Technical report, Stanford University, 2018.

Diana Farrell and Fiona Greig. Paychecks, paydays, and the online platform economy: Big data on income volatility. *JP Morgan Chase Institute*, 2016.

Jonathan V. Hall and Alan B. Krueger. An analysis of the labor market for uber driver-partners in the united states. Working Paper 22843, National Bureau of Economic Research, November 2016.

Etiennel Lale. Multiple jobholding over the past two decades. Technical report, Bureau of Labor Statistics, 2015.

Christina H Paxson and Nachum Sicherman. The dynamics of dual job holding and job mobility. *Journal of labor economics*, 14(3):357–393, 1996.

Francesco Renna and Ronald Oaxaca. The economics of dual job holding: A job portfolio model of labor supply. 2006.